

COMBINATORIAL SUBSTITUTIONS AND SOFIC TILINGS

THOMAS FERNIQUE AND NICOLAS OLLINGER

Laboratoire d'informatique fondamentale de Marseille (LIF)
Aix-Marseille Université, CNRS,
39 rue Joliot-Curie, 13 013 Marseille, France
E-mail address: {Thomas.Fernique,Nicolas.Ollinger}@lif.univ-mrs.fr

ABSTRACT. A combinatorial substitution is a map over tilings which allows to define sets of tilings with a strong hierarchical structure. In this paper, we show that such sets of tilings are sofic, that is, can be enforced by finitely many local constraints. This extends some similar previous results (Mozes'90, Goodman-Strauss'98) in a much shorter presentation.

1. Introduction

Tiling some space with geometrical shapes, or tiles, consists into covering this space with copies of the tiles. When a set of tilings can be characterized by adding finitely many local constraints on tiles so that the set of tilings corresponds exactly to the set of tilings satisfying the constraints, such a set of tilings is called sofic. Soficity corresponds to the interesting idea that the validity of a tiling can be locally proved by decorating tiles with constraints. In this paper, we contribute to a general question: which sets of tilings are sofic?

Substitutions provide a simple way to express how a set of tilings can be obtained by iteratively constructing bigger and bigger aggregates of tiles. The strong hierarchical structure of substitutions tilings permits to enforce global properties, for example aperiodicity. To prove that the set of tilings generated by a substitution is sofic, one has to encode the global hierarchical structure into local constraints on tiles. Such technique is at the root of classical papers on the undecidability of the Domino Problem [2, 9]. In the case of tilings on the square grid, Mozes [6] proved in a seminal paper that the set of tilings generated by rectangular non-deterministic substitutions satisfying a particular property is sofic. Goodman-Strauss [5] proved that such a construction method can be extended to a wide variety of geometrical substitutions. In this paper, based on ideas developed in [7], we further extend the construction to a broader class of substitutions by replacing the geometrical conditions by combinatorial conditions while decreasing the length of the presentation.

Let us sketch some definitions to state our main theorem. A sofic tiling is a valid tiling by a finite set of tiles with decorations. Combinatorial substitutions, introduced by Priebe-Frank in [8], map a tiling by tiles onto a tiling by so-called

Partially supported by the ANR projects EMC (ANR-09-BLAN-0164) and SubTile.

macro-tiles (finite tilings, here assumed to be connected), so that the macro-tiles of the latter are arranged as the tiles of the former. The limit set of a substitution is the set of complete tilings that admits preimages of any depth by the substitution. A good substitution is a substitution that is both connecting (a combinatorial condition ensuring that there is enough room for all information to flow) and consistent (a geometrical condition ensuring that the substitution can be correctly iterated). The main result obtained in this paper is:

Theorem 1.1. *The limit set of a good combinatorial substitution is sofic.*

The paper is organized as follows. Sections 2 and 3 formally define main notions, in particular sofic tilings and good combinatorial substitutions. Section 4 then presents *self-simulation*, which plays a central role in the constructive proof of Theorem 1.1, which is given in section 5. Last, section 6 concludes the paper by discussing an important parameter of this proof.

A single example illustrates definitions and results throughout the whole paper: the “Rauzy” example. It relies on the theory of *generalized substitutions* introduced in [1], a self-contained presentation of which is beyond the scope of this paper. Let us just mention that *Rauzy tilings* are digitizations of the planes of the Euclidean space with a specific given irrational normal vector. More details, as well as the results we here implicitly rely on, can be found in [3, 4]. We chose this example, not the simplest one, because it is not covered by results in [6, 5].

2. Sofic tilings

Polytopes are assumed to be homeomorphic to closed balls of \mathbb{R}^d and to have finitely many faces, with the $(d - 1)$ -dimensional faces being called *facets*.

A *tile* T is a polytope of the Euclidean space \mathbb{R}^d . A *tiling* Q of a *domain* $D \subset \mathbb{R}^d$ is a covering of D by interior-disjoint tiles, with the additional condition that two tiles can intersect (if they do) only along entire faces. A tiling is said to be *finite* if its domain is bounded.

Two tiles are said to be *adjacent* if they intersect along at least one facet, and a tiling is said to be *connected* if any two of its tiles can be connected by a sequence of adjacent tiles.

A facet of a tiling is said to be *external* if it is on the boundary of the domain, *internal* otherwise. We denote by ∂Q the set of external facets of a tiling Q . If this set is empty, then the tiling is said to be *complete*: its domain is the whole \mathbb{R}^d .

A *decorated tile* \mathcal{T} is a tile with a real map defined on its boundaries, called *decoration*. Two adjacent decorated tiles are said to *match* if their decorations are equal in any point of their intersecting facets. A *decorated tiling* \mathcal{Q} is a tiling by decorated tiles which pairwise match (when adjacent). Decorations are thus local constraints on the way tiles can be arranged in a tiling.

A *tileset* τ is a set of decorated tiles. It is said to be *finite* if it contains only a finite number of tiles up to direct isometries. A decorated tiling whose tiles belong to a tileset τ is called a τ -tiling; the set of τ -tilings is denoted by Λ_τ .

One says that a tiling *can be seen as* a decorated tiling if both are equal up to decorations. One denotes by π the map which removes the decorations. One easily

checks that any tiling can be seen as a τ -tiling if τ can be infinite. The interesting case is the one of finite tilesets:

Definition 2.1. A set of tilings is said to be *sofic* if it can be seen as the set of τ -tilings of some finite tileset τ .

For example, one can wonder whether the *Rauzy tilings* are sofic (Fig. 1).




Figure 1: A Rauzy tiling (partial view). Are Rauzy tilings sofic?

3. Combinatorial substitutions

Definition 3.1. A *combinatorial substitution* is a finite set of rules (P, Q, γ) , where P is a tile, Q is a finite connected tiling, and $\gamma : \partial P \rightarrow \partial Q$ maps distinct facets on disjoint sets of facets. The tiling Q is called a *macro-tile*, and if f is the k -th facet of P , then $\gamma(f)$ is called the k -th *macro-facet* of Q .

Fig. 2 illustrates this definition.

We call *tiling by macro-tiles* a tiling whose tiles can be partitioned into macro-tiles, with each macro-facet belonging to the intersection of exactly two macro-tiles. We can associate with each combinatorial substitution a binary relation over tilings:

Definition 3.2. Let σ be a combinatorial substitution. A tiling T by tiles of σ and a tiling T' by macro-tiles of σ are said to be σ -related if there is a one-to-one correspondence between the tiles of T and the macro-tiles of T' which preserves the combinatorial structure, that is, such that the a -th facet of a first tile of T matches the b -th facet of a second tile of T if and only if the a -th macro-facet of the first corresponding macro-tile of T' matches the b -th macro-facet of the second corresponding macro-tile of T' . One calls T a *preimage* of T' and T' an *image* of T .

For example, a Rauzy tiling can be uniquely seen as a tiling by Rauzy macro-tiles (Fig. 3) and has a unique preimage, which turns out to be itself a Rauzy tiling (both facts are non-trivial; they follow from results proven in [3]).

In particular, the relations associated with combinatorial substitutions yield a strong hierarchical structure on so-called limit-sets:




Figure 2: These five rules define the so-called Rauzy combinatorial substitution (a facet f and the corresponding macro-facet $\gamma(f)$ are similarly marked).




Figure 3: A Rauzy tiling can be uniquely seen as a tiling by Rauzy macro-tiles.

Definition 3.3. The *limit set* of a combinatorial substitution σ , denoted by Λ_σ , is the set of complete tilings which admit an infinite sequence of preimages.

For example, the limit set of the Rauzy combinatorial substitution is exactly the set of Rauzy tilings (again, this non-trivial fact follows from results proven in [3]).

Let us now turn to the *good combinatorial substitutions* to which Theorem 1.1 applies. First, a good combinatorial substitution must be *connecting*:

Definition 3.4. A combinatorial substitution σ is *connecting* if, for each rule (P, Q, γ) , the dual graph¹ of Q has a subgraph N , called its *network*, such that

- (1) N is a star with one branch for each macro-facet, and the leaf of the k -th branch is a tile with a facet, called k -th *port*, in the k -th macro-facet of Q ;

¹The dual graph of a tiling is the graph whose vertices correspond to tiles of the tiling and whose edges connect vertices corresponding to adjacent tiles.

- (2) Each macro-facet has non-port facets, and removing the edges of N and its central vertex yields a connected graph which connects² all these facets;
- (3) the center of N corresponds to a tile in the interior of Q , called *central tile*;
- (4) whenever two macro-tiles match along a port, they also match along the corresponding macro-facet.

Informally, the two first conditions ensure that the macro-facets are big enough to transfer via the network all the informations (encoded by decorations) that we need to enforce the hierarchical structure of the limit set. In particular, one easily sees that connectivity could not be achieved for combinatorial substitutions on the real line: Theorem 1.1 can apply only in dimension two or more. The third condition ensures that, by iteratively considering macro-tiles of macro-tiles (that is, when going higher and higher in the hierarchy of the limit set), we get tilings covering arbitrarily big balls. The last condition associates a port with its macro-facet (it is equivalent to the “sibling-edge-to-edge” condition of [5]).

Second, a good combinatorial substitution must be *consistent*:

Definition 3.5. A combinatorial substitution σ is said to be *consistent* if any tiling by macro-tiles of σ admits a preimage under σ .

Intuitively, consistency ensures that, if a tiling meets all the combinatorial conditions to have a preimage, then there is no geometrical obstruction to the existence of such a preimage.

Open Problem 3.6. Characterize consistent combinatorial substitutions.

For example, the Rauzy combinatorial substitution is connecting (one easily finds a suitable network for each rule, see Fig. 4 and 5) and consistent (this non-trivial fact follows from results in [4], where explicit maps defined over tilings are shown to be equivalent to such combinatorial substitutions). Theorem 1.1 thus yields that its limit set, *i.e.*, the set of Rauzy tilings, is sofic.




Figure 4: The Rauzy combinatorial substitution is connecting.

²One says that a subgraph *connects* a set of facets if these facets all belong to tiles which correspond to vertices of this subgraph.




Figure 5: A Rauzy macro-tile (top-left) and an image of it under the Rauzy combinatorial substitution, that one could call Rauzy “macro-macro-tile” (center).

4. Self-simulation

Definition 4.1. Let σ be a combinatorial substitution with the rules $\{(P_i, Q_i, \gamma_i)\}_i$. A tileset τ is said to σ -self-simulates if there is a set of τ -tilings, called τ -macro-tiles, and a map ϕ from these τ -macro-tiles into τ such that

- (1) for any τ -macro-tile \mathcal{Q} , there is i such that $\pi(\mathcal{Q}) = Q_i$ and $\pi(\phi(\mathcal{Q})) = P_i$;
- (2) any complete τ -tiling is also a tiling by τ -macro-tiles;
- (3) the a -th macro-facet of a τ -macro-tile \mathcal{Q} can match the b -th macro-facet of a τ -macro-tile \mathcal{Q}' if and only if the a -th facet of the τ -tile $\phi(\mathcal{Q})$ can match the b -th facet of the τ -tile $\phi(\mathcal{Q}')$.

Proposition 4.2. *If a tileset σ -self-simulates for a consistent combinatorial substitution σ , then its complete tilings are, up to decorations, in the limit set of σ .*

Proof. Consider a complete τ -tiling \mathcal{P} . Conditions (1)–(2) ensure that removing the decorations of \mathcal{P} yields a tiling by macro-tiles of σ , say P . The consistency of σ ensures that P admits a preimage under σ , say R . Let us show that R can be endowed by decorations to get a τ -tiling. Consider a tile T of R . This tile corresponds (via the one-to-one correspondence in Def. 3.2) to a macro-tile Q of P , which is itself the image under π of a τ -macro-tile \mathcal{Q} of \mathcal{P} . We associate with T the τ -tile $\phi(\mathcal{Q})$. Condition (3) ensures that replacing tiles in R by their such associated τ -tiles yields a complete τ -tiling, say \mathcal{R} . We can repeat all this process, with \mathcal{R} instead of \mathcal{P} . By induction, we get an infinite sequence of complete tilings, with each one being the preimage under σ of the previous one. Thus, $\mathcal{P} \in \Lambda_\sigma$. ■

5. Constructive proof of Theorem 1.1

Let σ be a good combinatorial substitution with rules $(P_i, Q_i, \gamma_i)_i$. We here rely on the fact that σ is connecting to construct a finite tileset τ which σ -self-simulates. The consistency of σ then ensures, via Prop. 4.2, that $\pi(\Lambda_\tau) \subseteq \Lambda_\sigma$ holds. We also show that the converse inclusion holds. This thus constructively proves Theorem 1.1.

5.1. Settings

Let T_1, \dots, T_n and f_1, \dots, f_m be numberings of, respectively, the tiles and the internal facets of all the Q_i 's. Given the k -th facet of the tile T_i , $N_\sigma(i, k)$ stands either for its index if it is an internal facet, or for a special value “port”, “macro-facet” or “boundary” otherwise (depending whether it is a port, a non-port facet in a macro-facet or another external facet).

Each tile of τ is a T_i endowed with a decoration which encodes on each facet a triple (f, j, g) , where f and g are either facet indices or special values “port”, “macro-facet” or “boundary”, and j is either zero or a tile index. We call f the *macro-index*, j the *parent-index* and g the *neighbor-index*. This clearly allows only a finite number of different tiles.

We skip the technical details concerning the way these triples are encoded by decorations³. We assume that two decorated tiles match along a facet if and only if the same triple is encoded on both facets, and that the only direct isometry which leaves invariant the decoration of a facet is the identity, so that a decorated tile cannot trivially match with a translated or rotated copy of itself.

³Recall that the decoration of a tile is a real map defined on its boundary.

5.2. Decorations

The five following steps completely define a tiling τ .

1. The macro-index of the k -th facet of any decorated T_j is $N_\sigma(j, k)$. This step ensures that any complete τ -tiling can be uniquely seen as a tiling by τ -macro-tiles.
2. Consider a decorated non-central tile of Q_i . Its facets which are internal and not crossed by the network have all the same parent-index, also called parent-index of the tile, which can be any j such that $T_j = P_i$. Its facets which are external, port excluded, have parent-index 0. This step ensures that the τ -tiles of a τ -macro-tile Q share a common parent-index j ; the tile T_j is called the *parent-tile* of Q .
3. Consider, in a non-central τ -tile with parent-index j , a facet which is neither a port nor crossed by the network. Its neighbor-index is either $N_\sigma(j, k)$ if it is in the k -th macro-facet, or equal to its macro-index otherwise. This step ensures that the macro-facets (ports excepted) of a τ -macro-tile are equivalent to the macro-indices of its parent tile (once decorated).
4. Consider a non-central τ -tile with parent-index j . Its facets which are either k -th port or crossed by the k -th branch of the network have all the same pair of parent/neighbor indices: it can be any pair not forbidden on the k -th facet of T_j (Steps 1–3). This step ensures that each port of a τ -macro-tile is equivalent to the pair of parent/neighbor indices of a decorated parent-tile⁴.
5. Whenever a non-central τ -tile \mathcal{T} has as many facets as a central tile T_j , we define a central τ -tile \mathcal{T}' by endowing each k -th facet of T_j with the pair of parent/neighbor indices on the k -th facet of \mathcal{T} (the macro-indices are defined as usual, see Step 1). One says that \mathcal{T}' *derives* from \mathcal{T} .

5.3. First inclusion

Let us show that the above defined tiling σ -self-simulates. Given a τ -macro-tile Q with parent-index j and central τ -tile \mathcal{T}' , let $\phi(Q)$ be the decorated tile obtained by endowing the k -th facet of T_j with the parent/neighbor indices on the k -th facet of \mathcal{T}' (the macro-indices are defined as for any tile, see Step 1). One checks that ϕ is a map satisfying conditions (1)–(3) of Def. 4.1. It remains to check that $\phi(Q) \in \tau$.

Let \mathcal{T} be a non-central τ -tile from which derives \mathcal{T}' . If T_j is central, then $\phi(Q)$ also derives from \mathcal{T} , hence is in τ . Otherwise, Step 4 ensures, for each k , that the parent/neighbor indices on the k -th facet of \mathcal{T}' appear on the k -th facet of a decorated T_j . In particular, this holds on facets of T_j which are not crossed by a network. Since the neighbor-indices of these facets can only⁵ appear on a T_j (see Steps 1 and 3), \mathcal{T} is a decorated T_j . This yields $\phi(Q) = \mathcal{T}$, and thus $\phi(Q)$ is in τ .

Prop. 4.2 then applies and yields the first inclusion $\pi(\Lambda_\tau) \subseteq \Lambda_\sigma$.

⁴With the problem that a parent-tile can be decorated in different ways, and nothing yet prevents the ports from mixing these decorations.

⁵Actually, a facet-index appears on the two tiles of a macro-tile which share the corresponding facet. We thus need, in order to completely characterize T_j , either to assume that there is at least two such facets (this is a rather mild assumption), or to endow facets with an *orientation* and to allow two facets to match if and only if they have opposite orientations.

5.4. Second inclusion

Let us extend τ in a tileset τ' as follow. For each τ -tile \mathcal{T} and each subset S of its facets, we define a τ' -tile \mathcal{T}_S by replacing the decorations of the facets in S by a special decoration “undefined”.

Now, let P be a tiling in Λ_σ . Consider an infinite sequence $(P_n)_{n \geq 0}$ of successive preimages of P . Given $n > 0$, P_n can be seen as a τ' -tiling: it suffices to endow any facet with “undefined”. Then, P_{n-1} can be seen as a tiling by τ' -macro-tiles, with “undefined” decorations appearing only on the network. Indeed, consider a macro-tile of P_{n-1} which corresponds (via the one-to-one correspondence in Def. 3.2) to a tile T_j in P_n : it suffices to endow its tiles as in steps 1–3 of the definition of τ , with T_j being the parent-tile, and with the decoration “undefined” on the facets crossed by the network. This can be iterated up to $P = P_0$, and the third condition of Def. 3.4 ensures that the decorations “undefined” appear only on sort of grids whose cells have bigger and bigger size.

Thus, by making n tend to infinity, one can see P as a τ' -tiling whose “undefined” decorations, if any, form either a star with k infinite branches, or a single biinfinite branch. In the first case, we can replace the central τ' -tile by any τ -tile with k facets, and then the tiles on branches by τ -tiles which carry decorations to infinity. In the second case, we can replace the τ' -tiles by τ -tiles which carry any decoration on the whole branch. In any case, we can thus see P as a τ -tiling.

We thus have the second inclusion $\Lambda_\sigma \subseteq \pi(\Lambda_\tau)$. Both inclusions prove Theorem 1.1.

6. On the number of tiles

Let us conclude this paper by discussing the size $\#\tau$ of the tileset τ defined in the previous section (although its finiteness suffices for Theorem 1.1).

Consider a good combinatorial substitution with r rules. Fix a network as in Def. 3.4 for each of these rules. Let n and m denote the total number of, respectively, tiles and internal facets of the macro-tiles of these r rules. Among these n tiles, let p denote the number of those which lie on a network.

First, the $n - p$ tiles not on the network can be decorated in at most n different ways, according to the parent-index they carry. This yields at most $N_0 = (n - p)n$ τ -tiles. Then, each of the $p - r$ non-central tiles on the network can be decorated in at most $(n - p)n + pnm$ different ways: $n - p$ parent-indices which correspond to tiles not on the network, hence allow at most n pairs of parent/neighbor indices carried on the network, and p parent-indices which correspond to tiles on the network, hence allow at most nm pairs of parent/neighbor indices carried on the network. This yields at most $N_p = (p - r)((n - p)n + pmn)$ τ -tiles. Last, the r central tiles can be decorated in at most as many different ways as there is non-central τ -tiles. Finally, this yields

$$\#\tau \leq (r + 1)(N_0 + N_p) \leq (r + 2)p^2mn.$$

This bound is, for example, about one billion in the Rauzy case.

In order to reduce this huge number of tiles, it is worth noting that, instead of carrying a parent-index through all the tiles of a macro-tile, it suffices to carry




Figure 6: A second network for the Rauzy combinatorial substitution.

it along a *second network* connecting the macro-facets of this macro-tile and intersecting each of the branches of its (first) network (see, *e.g.*, Fig. 6). The “control” described in Step 4 is then performed only on tiles where both networks crosses, while we simply allow any possible pairs of parent/neighbor indices to be carried through the tiles which are only on the first network. If we denote by q the total number of tiles on these new second networks and by c be the total number of crossings between second and first networks, a similar analysis yields

$$\#\tau \leq (r+1)(N'_0 + N'_q + N'_p + N'_c) \leq (r+2)cp(m+qn),$$

where:

$$\begin{aligned} N'_0 &= n - p - q + c, & N'_p &= (p - c)(m + qn), \\ N'_q &= (q - c)n, & N'_c &= c(N'_0 + N'_q + N'_p). \end{aligned}$$

This bound is, for example, about 70 millions tiles in the Rauzy case.

This last bound is huge but generic. One can hope to dramatically decrease this bound in specific cases. Indeed, most of the τ -tiles correspond to tiles which simply carry any possible information (the tiles on networks, crossings excepted). Since these tiles all play the same role, it would be worth to replace all the tiles on a network by a single tile (one can thus hope to gain a factor pq in the above bound – this would yield about 25000 tiles in the Rauzy case). This shall however be done carefully, so that tiles still necessarily form macro-tiles. Note that it should be much easier in dimension $d \geq 3$, since the cohesion of macro-tiles can be more easily enforced without relying on tiles on networks.

Acknowledgments. We would like to thank the anonymous referees for their valuable comments. The first author also would like to thank N. Pytheas Fogg and C. Goodman-Strauss for their encouragements to write down this result, as well as for useful discussions.

References

- [1] P. Arnoux and S. Ito, *Pisot substitutions and Rauzy fractals*, Bull. Belg. Math. Soc. Simon Stevin **8** (2001), no. 2, 181–207.
- [2] R. Berger, *The undecidability of the domino problem*, Ph.D. thesis, Harvard University, July 1964.
- [3] V. Berthé and Th. Fernique, *Brun expansions of stepped surfaces*, Disc. Math. **311** (2011), no. 7, 521–543.
- [4] Th. Fernique, *Local rule substitutions and stepped surfaces*, Theor. Comput. Sci. **380** (2007), no. 3, 317–329.
- [5] C. Goodman-Strauss, *Matching rules and substitution tilings*, Ann. of Math. (2) **147** (1998), no. 1, 181–223.
- [6] S. Mozes, *Tilings, substitution systems and dynamical systems generated by them*, J. Analyse Math. **53** (1989), 139–186.
- [7] N. Ollinger, *Two-by-two substitution systems and the undecidability of the domino problem*, Proceedings of CiE'2008, LNCS, vol. 5028, Springer, 2008, pp. 476–485.
- [8] N. Priebe Frank, *Detecting combinatorial hierarchy in tilings using derived voronoi tessellations*, Disc. Comput. Geom. **29** (2003), no. 3, 459–467.
- [9] R. M. Robinson, *Undecidability and nonperiodicity for tilings of the plane*, Inventiones Mathematicae **12** (1971), 177–209.

Appendix A. Complete example: the 3×3 square substitution




Figure 7: a simple substitution



Figure 8: sample σ -related tilings

Figure 9: chosen network and numberings




Figure 10: decoration of facets




Figure 11: initial decoration scheme




Figure 12: after step 1: macro-indices are fixed






Figure 13: after step 2: parent-indices are fixed outside network




where $n = N_\sigma(j, N)$, $s = N_\sigma(j, S)$, $w = N_\sigma(j, W)$ and $e = N_\sigma(j, E)$.

Figure 14: after step 3: all decorations fixed outside network



where $w = N_\sigma(j, W)$, for any j and any horizontal pair (α, β) of parent/neighbor indices

Figure 15: after step 4: all decorations fixed but for central tiles



where, for any j , $n = N_\sigma(j, N)$, $s = N_\sigma(j, S)$, $w = N_\sigma(j, W)$, $e = N_\sigma(j, E)$, (α_j^k, β_j^k) are valid parent/neighbor indices for a tile T_k with parent-index j

Figure 16: after step 5: all decorations are fixed